Everything old is new again: Quoted Domain Specific Languages

Shayan Najd Sam Lindley
The University of Edinburgh
{sh.najd,sam.lindley}@ed.ac.uk

Josef Svenningsson
Chalmers University of Technology
josefs@chalmers.se

Philip Wadler
The University of Edinburgh
wadler@inf.ed.ac.uk

arXiv:1507.07264v2 [cs.PL] 4 Aug 2015

Abstract

We describe a new approach to domain specific languages (DSLs), called Quoted DSLs (QDSLs), that resurrects two old ideas: quotation, from McCarthy's Lisp of 1960, and the subformula property, from Gentzen's natural deduction of 1935. Quoted terms allow the DSL to share the syntax and type system of the host language. Normalising quoted terms ensures the subformula property, which guarantees that one can use higher-order types in the source while guaranteeing first-order types in the target, and enables using types to guide fusion. We test our ideas by re-implementing Feldspar, which was originally implemented as an Embedded DSL (EDSL), as a QDSL; and we compare the QDSL and EDSL variants.

Categories and Subject Descriptors D.1.1 [*Applicative (Functional) Programming*]; D.3.1 [*Formal Definitions and Theory*]; D.3.2 [*Language Classifications*]: Applicative (functional) languages

Keywords domain-specific language, DSL, EDSL, QDSL, embedded language, quotation, normalisation, subformula property

1. Introduction

> Don't throw the past away
> You might need it some rainy day
> Dreams can come true again
> When everything old is new again
>
> – Peter Allen and Carole Sager

Implementing domain-specific languages (DSLs) via quotation is one of the oldest ideas in computing, going back at least to McCarthy's Lisp, which was introduced in 1960 and had macros as early as 1963. Today, a more fashionable technique is Embdedded DSLs (EDSLs), which may use shallow embedding, deep embedding, or a combination of the two. Our goal in this paper is to reinvigorate the idea of building DSLs via quotation, by introducing an approach we dub Quoted DSLs (QDSLs). A key feature of QDSLs is the use of normalisation to ensure the subformula property, first proposed by Gentzen in 1935.

[Copyright notice will appear here once 'preprint' option is removed.]

> Imitation is the sincerest of flattery.
>
> — Charles Caleb Colton

Cheney et al. (2013) describe a DSL for language-integrated query in F# that translates into SQL. Their technique depends on quotation, normalisation of quoted terms, and the subformula property—an approach which we here dub QDSL. They conjecture that other DSLs might benefit from the same technique, particularly those that perform staged computation, where host code at generation-time computes target code to be executed at run-time.

Generality starts at two. Here we test the conjecture of Cheney et al. (2013) by reimplementing the EDSL Feldspar (Axelsson et al. 2010) as a QDSL. We describe the key features of the design, and show that the performance of the two versions is comparable. We compare the QDSL and EDSL variants of Feldspar, and assess the tradeoffs between the two approaches.

Davies and Pfenning (2001) also suggest quotation as a foundation for staged computation, and note a propositions-as-types connection between quotation and a modal logic; our type $Qt\ a$ corresponds to their type $\bigcirc a$. They also mention in passing the utility of normalising quoted terms, although they do not mention the subformula property.

The .NET Language-Integrated Query (LINQ) framework as used in C# and F# (Meijer et al. 2006; Syme 2006), and the Lightweight Modular Staging (LMS) framework as used in Scala (Rompf and Odersky 2010), exhibit overlap with the techniques described here. Notably, they use quotation to represent staged DSL programs, and they make use to a greater or lesser extent of normalisation. In F# LINQ quotation is indicated in the normal way (by writing quoted programs inside special symbols), while in C# LINQ and Scala LMS quotation is indicated by type inference (quoted terms are given a special type).

> Perhaps we may express the essential properties of such a normal proof by saying: it is not roundabout.
>
> — Gerhard Gentzen

Our approach exploits the fact that normalised terms satisfy the subformula property, first introduced in the context of natural deduction by Gentzen (1935), and improved by Prawitz (1965).

The subformulas of a formula are its subparts; for instance, the subformulas of $A \to B$ are the formula itself and the subformulas of A and B. The subformula property states that every proof can be put into a normal form where the only propositions that appear in the proof are subformulas of the hypotheses and conclusion of the proof. Applying the principle of Propositions as Types (Howard 1980; Wadler 2015), the subformula property states that every lambda term can be put into a normal form where the only types that appear in the term are subformulas of the types of the free variables and the type of the term itself.

The subformula property provides users of the DSL with useful guarantees, such as the following:

- they may write higher-order terms while guaranteeing to generate first-order code;
- they may write a sequence of loops over arrays while guaranteeing to generate code that fuses those loops;
- they may write intermediate terms with nested collections while guaranteeing to generate code that operates on flat data.

The first and second are used in this paper, and are key to generating C; while the first and third are used by Cheney et al. (2013) and are key to generating SQL.

The subformula property is closely related to conservativity. A conservativity result expresses that adding a feature to a system of logic, or to a programming language, does not make it more expressive. Consider intuitionistic logic with conjunction; conservativity states that adding implication to this logic proves no additional theorems that can be stated in the original logic. Such a conservativity result is an immediate consequence of the subformula property; since the hypotheses and conjuction of the proof only mention conjunction, any proof, even if it uses implication, can be put into a normal form that only uses conjunction. Equivalently, any lambda calculus term that mentions only pair types in its free variables and result, even if it uses functions, can be put in a normal form that only uses pairs. Such a result is related to the first bullet point above; see Proposition 4.4 in Section 4.

As another example, the third bullet point above corresponds to a standard conservativity result for databases, namely that nested queries are no more expressive than flat queries (Wong 1996). This conservativity result, as implied by the subformula property, is used by Cheney et al. (2013) to show that queries that use intermediate nesting can be translated to SQL, which only queries flat tables and does not support nesting of data.

The subformula property holds only for terms in normal form. Previous work, such as (Cheney et al. 2013) uses call-by-name normalisation algorithm that may cause computations to be repeated. Here we present call-by-value and call-by-need normalisation algorithms that guarantee to preserve sharing of computations. We also present a sharpened version of the subformula property, which we apply to characterise the circumstances under which a QDSL may guarantee to generate first-order code.

> Good artists copy, great artists steal. — Picasso

EDSL is great in part because it steals the type system of its host language. Arguably, QDSL is greater because it steals the type system, the syntax, and the normalisation rules of its host language.

In theory, an EDSL should also steal the syntax of its host language, but in practice the theft is often only partial. For instance, an EDSL such as Feldspar (Axelsson et al. 2010) or Nikola (Mainland and Morrisett 2010), when embedded in Haskell, can exploit overloading so that arithmetic operations in both languages appear identical, but the same is not true of comparison or conditionals. In QDSL, the syntax of the host and embedded languages is identical. For instance, this paper presents a QDSL variant of Feldspar, again in Haskell, where arithmetic, comparison, and conditionals are all represented by quoted terms, and hence identical to the host.

An EDSL may also steal the normalisation rules of its host language, using evaluation in the host to normalise terms of the target, but again the theft is often only partial. Section 5 compares QDSL and EDSL variants of Feldspar. In the first example, it is indeed the case that the EDSL achieves by evaluation of host terms what the QDSL achieves by normalisation of quoted terms. However, in other cases, the EDSL must perform normalisation of the deep embedding corresponding to how the QDSL normalises quoted terms.

> Try to give all of the information to help others to judge the value of your contribution; not just the information that leads to judgment in one particular direction or another.
> — Richard Feynman

The subformula property depends on normalisation, but normalisation may lead to exponential blowup in the size of the normalised code when there are nested conditionals; and hyperexponential blowup in recondite cases involving higher-order functions. We explain how uninterpreted constants allow the user to control where normalisation does and does not occur, while still maintaining the subformula property. Future work is required to consider trade-offs between full normalisation as required for the subformula property and special-purpose normalisation as used in many DSLs; possibly a combination of both will prove fruitful.

Some researchers contend an essential property of a DSL which generates target code is that every type-correct term should successfully generate code in the target language. EDSL Feldspar satisfies this property; but neither P-LINQ of Cheney et al. (2013) nor QDSL Feldspar satisfy this property, since the user is required to eyeball quoted code to ensure it mentions only permitted operators. If this is thought too onerous, it is possible to ensure the property with additional preprocessing.

> This is the short and the long of it. — Shakespeare

The contributions of this paper are:

- To introduce QDSLs as an approach to building DSLs based on quotation, normalisation of quoted terms, and the subformula property by presenting the design of a QDSL variant of Feldspar (Section 2).
- To measure QDSL and EDSL implementations of Feldspar, and show they offer comparable performance (Section 3).
- To present normalisation algorithms for call-by-value and call-by-need that preserve sharing, and to formulate a sharpened version of the subformula property and apply it to characterise when higher-order terms normalise to first-order form (Section 4).
- To compare the QDSL variant of Feldspar with the deep and shallow embedding approach used in the EDSL variant of Feldspar, and show they offer tradeoffs with regard to ease of use (Section 5).

Section 6 describes related work, and Section 7 concludes.

Our QDSL and EDSL variants of Feldspar and benchmarks are available at `https://github.com/shayan-najd/QFeldspar`.

2. Feldspar as a QDSL

Feldspar is an EDSL for writing signal-processing software, that generates code in C (Axelsson et al. 2010). We present a variant, QDSL Feldspar, that follows the structure of the previous design closely, but using the methods of QDSL rather than EDSL. Section 5 compares the QDSL and EDSL designs.

2.1 The top level

In QDSL Feldspar, our goal is to translate a quoted term to C code. The top-level function has the type:

$$qdsl :: (Rep\ a, Rep\ b) \Rightarrow Qt\ (a \rightarrow b) \rightarrow C$$

Here $Qt\ \ a$ represents a Haskell term of type a, its *quoted* representation, and type C represents code in C. The top-level function

expects a quoted term representing a function from type a to type b, and returns C code that computes the function.

Not all types representable in Haskell are easily representable in C. For instance, we do not wish our target C code to manipulate higher-order functions. The argument type a and result type b of the main function must be representable, which is indicated by the type-class restrictions $Rep\ a$ and $Rep\ b$. Representable types include integers, floats, and pairs where the components are both representable.

$$\textbf{instance}\ Rep\ Int$$
$$\textbf{instance}\ Rep\ Float$$
$$\textbf{instance}\ (Rep\ a, Rep\ b) \Rightarrow Rep\ (a, b)$$

2.2 A first example

Let's begin with the "hello world" of program generation, the power function. Since division by zero is undefined, we arbitrarily choose that raising zero to a negative power yields zero. Here is an optimised power function represented using QDSL:

$$\begin{array}{l} power :: Int \to Qt\ (Float \to Float) \\ power\ n = \\ \quad \textbf{if}\ n < 0\ \textbf{then} \\ \quad\quad [||\lambda x \to \textbf{if}\ x == 0\ \textbf{then}\ 0 \\ \quad\quad\quad\quad\quad\quad \textbf{else}\ \ 1\ /\ (\$\$(power\ (-n))\ x)||] \\ \quad \textbf{else if}\ n == 0\ \textbf{then} \\ \quad\quad [||\lambda x \to 1||] \\ \quad \textbf{else if}\ even\ n\ \textbf{then} \\ \quad\quad [||\lambda x \to \textbf{let}\ y = \$\$(power\ (n\ \mathrm{div}\ 2))\ x\ \textbf{in}\ y \times y||] \\ \quad \textbf{else} \\ \quad\quad [||\lambda x \to x \times (\$\$(power\ (n-1))\ x)||] \end{array}$$

The typed quasi-quoting mechanism of Template Haskell is used to indicate which code executes at which time. Unquoted code executes at generation-time while quoted code executes at run-time. Quoting is indicated by $[||...||]$ and unquoting by $\$\$(\cdots)$.

Invoking $qdsl\ (power\ (-6))$ generates code to raise a number to its -6 power. Evaluating $power\ (-6)$ yields the following:

$$\begin{array}{l} [||\lambda x \to \textbf{if}\ x == 0\ \textbf{then}\ 0\ \textbf{else}\ 1\ / \\ \quad (\lambda x \to \textbf{let}\ \{y = (\lambda x \to x \times \\ \quad\quad (\lambda x \to \textbf{let}\ \{y = (\lambda x \to x \times (\lambda x \to 1)\ x)\ x\} \\ \quad\quad\quad \textbf{in}\ y \times y)\ x)\ x\}\ \textbf{in}\ y \times y)\ x||] \end{array}$$

Normalising as described in Section 4, with variables renamed for readability, yields the following:

$$\begin{array}{l} [||\lambda u \to \textbf{if}\ u == 0\ \textbf{then}\ 0\ \textbf{else} \\ \quad \textbf{let}\ v = u \times 1\ \textbf{in} \\ \quad \textbf{let}\ w = u \times (v \times v)\ \textbf{in} \\ \quad 1\ /\ (w \times w)||] \end{array}$$

With the exception of the top-level term, all of the overhead of lambda abstraction and function application has been removed; we explain below why this is guaranteed by the subformula property. From the normalised term it is easy to generate the final C code:

```
float prog (float u) {
  float w; float v; float r;
  if (u == 0.0) {
    r = 0.0;
  } else {
    v = (u * 1.0);
    w = (u * (v * v));
    r = (1.0f / (w * w));
  }
  return r;
}
```

By default, we always generate a routine called `prog`; it is easy to provide the name as an additional parameter if required.

Depending on your point of view, quotation in this form of QDSL is either desirable, because it makes manifest the staging, or undesirable because it is too noisy. QDSL enables us to "steal" the entire syntax of the host language for our DSL. In Haskell, an EDSL can use the same syntax for arithmetic operators, but must use a different syntax for equality tests and conditionals, as explained in Section 5.

Within the quotation brackets there appear lambda abstractions and function applications, while our intention is to generate first-order code. How can the QDSL Feldspar user be certain that such function applications do not render transformation to first-order code impossible or introduce additional runtime overhead? The answer is the subformula property.

2.3 The subformula property

Gentzen's subformula property guarantees that any proof can be normalised so that the only formulas that appear within it are subformulas of one of the hypotheses or of the conclusion of the proof. Viewed through the lens of Propositions as Types, Gentzen's subformula property guarantees that any term can be normalised so that the type of each of its subterms is a subformula of either the type of one of its free variables (corresponding to hypotheses) or of the term itself (corresponding to the conclusion). Here the subformulas of a type are the type itself and the subformulas of its parts, where the parts of $a \to b$ are a and b, the parts of (a, b) are a and b, and types Int and $Float$ have no parts. (See Proposition 4.2.)

Further, it is easy to adapt the original proof to guarantee a sharpened subformula property: any term can be normalised so that the type of each of its proper subterms is a proper subformula of either the type of one of its free variables (corresponding to hypotheses) or the term itself (corresponding to the conclusion). Here the proper subterms of a term are all subterms save for free variables and the term itself, and the proper subformulas of a type are all subformulas save for the type itself. In the example of the previous subsection, the sharpened subformula property guarantees that after normalisation a closed term of type $float \to float$ will only have proper subterms of type $float$, which is indeed true for the normalised term. (See Proposition 4.3.)

The subformula property depends on normalisation of terms, but complete normalisation is not always possible or desirable. The extent of normalisation may be controlled by introducing uninterpreted constants. In particular, we introduce the uninterpreted constant

$$save :: Rep\ a \Rightarrow a \to a$$

of arity 1, which is equivalent to the identity function on representable types. Unfolding of an application $L\ M$ can be inhibited by rewriting it in the form $save\ L\ M$, where L and M are arbitrary terms. A use of $save$ appears in Section 2.6. In a context with recursion, we take

$$fix :: (a \to a) \to a$$

as an uninterpreted constant.

2.4 A second example

In the previous code, we arbitrarily chose that raising zero to a negative power yields zero. Say that we wish to exploit the $Maybe$ type of Haskell to refactor the code, by separating identification of the exceptional case (negative exponent of zero) from choosing a value for this case (zero). We decompose $power$ into two functions $power'$ and $power''$, where the first returns $Nothing$ in the exceptional case, and the second maps $Nothing$ to a suitable value.

The $Maybe$ type is a part of the Haskell standard prelude.

```
data Maybe a = Nothing | Just a
maybe :: b → (a → b) → Maybe a → b
return :: a → Maybe a
(>>=)  :: Maybe a → (a → Maybe b) → Maybe b
```

Here is the refactored code.

```
power' :: Int → Qt (Float → Maybe Float)
power' n =
  if n < 0 then
    [||λx → if x == 0 then Nothing
                      else do y ← $$(power' (−n)) x
                              return (1 / y)||]
  else if n == 0 then
    [||λx → return 1||]
  else if even n then
    [||λx → do y ← $$(power' (n div 2)) x
              return (y × y)||]
  else
    [||λx → do y ← $$(power' (n − 1)) x
              return (x × y)||]
power'' :: Int → Qt (Float → Float)
power'' n =
  [||λx → maybe 0 (λy → y) ($$(power' n) x)||]
```

Evaluation and normalisation of $power\ (-6)$ and $power''\ (-6)$ yield identical terms (up to renaming), and hence applying *qdsl* to these yields identical C code.

The subformula property is key: because the final type of the result does not involve *Maybe* it is certain that normalisation will remove all its occurrences. Occurrences of **do** notation are expanded to applications of $(\gg=)$, as usual. Rather than taking *return*, $(\gg=)$, and *maybe* as uninterpreted constants (whose types have subformulas involving *Maybe*), we treat them as known definitions to be eliminated by the normaliser. Type *Maybe* is a sum type, and is normalised as described in Section 4.

2.5 While

Code that is intended to compile to a `while` loop in C is indicated in QDSL Feldspar by application of *while*.

```
while :: (Rep s) ⇒ (s → Bool) → (s → s) → s → s
```

Rather than using side-effects, *while* takes three arguments: a predicate over the current state, of type $s \to Bool$; a function from current state to new state, of type $s \to s$; and an initial state of type s; and it returns a final state of type s. So that we may compile *while* loops to C, the type of the state is constrained to representable types.

We can define a *for* loop in terms of a *while* loop.

```
for :: (Rep s) ⇒ Qt (Int → s → (Int → s → s) → s)
for = [||λn s₀ b → snd (while (λ(i, s) → i < n)
                              (λ(i, s) → (i + 1, b i s))
                              (0, s₀))||]
```

The state of the *while* loop is a pair consisting of a counter and the state of the *for* loop. The body b of the *for* loop is a function that expects both the counter and the state of the *for* loop. The counter is discarded when the loop is complete, and the final state of the *for* loop returned. Here *while*, like *snd* and $(+)$, is a constant known to QDSL Feldspar, and so not enclosed in $$ antiquotes.

As an example, we can define Fibonacci using a *for* loop.

```
fib :: Qt (Int → Int)
fib = [||λn → fst ($$for n (0, 1) (λi (a, b) → (b, a + b)))||]
```

Again, the subformula property plays a key role. As explained in Section 2.3, primitives of the language to be compiled, such as $(\times)$ and *while*, are treated as free variables or constants of a given arity. As described in Section 4, we can ensure that after normalisation every occurence of *while* has the form

$$while\ (\lambda s \to \cdots)\ (\lambda s \to \cdots)\ (\cdots)$$

where the first ellipses has type *Bool*, and both occurrences of the bound variable s and the second and third ellipses all have the same type, that of the state of the while loop.

Unsurprisingly, and in accord with the subformula property, each occurrence of *while* in the normalised code will contain subterms with the type of its state. The restriction of state to representable types increases the utility of the subformula property. For instance, since we have chosen that *Maybe* is not a representable type, we can ensure that any top-level function without *Maybe* in its type will normalise to code not containing *Maybe* in the type of any subterm. In particular, *Maybe* cannot appear in the state of a *while* loop, which is restricted to representable types. An alternative choice is possible, as we will see in the next section.

2.6 Arrays

A key feature of Feldspar is its distinction between two types of arrays, manifest arrays, *Arr*, which may appear at run-time, and "pull arrays", *Vec*, which are eliminated by fusion at generation-time. Again, we exploit the subformula property to ensure no subterms of type *Vec* remain in the final program.

The type *Arr* of manifest arrays is simply Haskell's array type, specialised to arrays with integer indices and zero-based indexing. The type *Vec* of pull arrays is defined in terms of existing types, as a pair consisting of the length of the array and a function that given an index returns the array element at that index.

```
type Arr a = Array Int a

data Vec a = Vec Int (Int → a)
```

Values of type *Arr* are representable, assuming that the element type is representable, while values of type *Vec* are not representable.

```
instance (Rep a) ⇒ Rep (Arr a)
```

For arrays, we assume the following primitive operations.

```
mkArr :: (Rep a) ⇒ Int → (Int → a) → Arr a
lnArr :: (Rep a) ⇒ Arr a → Int
ixArr :: (Rep a) ⇒ Arr a → Int → a
```

The first populates a manifest array of the given size using the given indexing function, the second returns the length of the array, and the third returns the array element at the given index. Array components must be representable.

We define functions to convert between the two representations in the obvious way.

```
toVec   :: Rep a ⇒ Qt (Arr a → Vec a)
toVec   = [||λa → Vec (lnArr a) (λi → ixArr a i)||]
fromVec :: Rep a ⇒ Qt (Vec a → Arr a)
fromVec = [||λ(Vec n g) → mkArr n g||]
```

It is straightforward to define operations on vectors, including combining corresponding elements of two vectors, summing the elements of a vector, dot product of two vectors, and norm of a vector. When combining two vectors, the length of the result is the minimum of the lengths of the arguments.

```
minim  :: Ord a ⇒ Qt (a → a → a)
minim  = [||λx y → if x < y then x else y||]
```

$$
\begin{array}{ll}
zipVec & :: Qt\ ((a \to b \to c) \to Vec\ a \to Vec\ b \to Vec\ c) \\
zipVec & = [||\lambda f\ (Vec\ m\ g)\ (Vec\ n\ h) \to \\
 & \qquad Vec\ (\$\$minim\ m\ n)\ (\lambda i \to f\ (g\ i)\ (h\ i))||] \\
sumVec & :: (Rep\ a, Num\ a) \Rightarrow Qt\ (Vec\ a \to a) \\
sumVec & = [||\lambda(Vec\ n\ g) \to \$\$for\ n\ 0\ (\lambda i\ x \to x + g\ i)||] \\
dotVec & :: (Rep\ a, Num\ a) \Rightarrow Qt\ (Vec\ a \to Vec\ a \to a) \\
dotVec & = [||\lambda u\ v \to \$\$sumVec\ (\$\$zipVec\ (\times)\ u\ v)||] \\
normVec & :: Qt\ (Vec\ Float \to Float) \\
normVec & = [||\lambda v \to sqrt\ (\$\$dotVec\ v\ v)||]
\end{array}
$$

The third of these uses the *for* loop defined in Section 2.5.

Our final function cannot accept Vec as input, since the Vec type is not representable, but it can accept Arr as input. For instance, if we invoke *qdsl* on

$$[||\ \$\$\ normVec \circ \$\$toVec||]$$

the quoted term normalises to

$$
\begin{array}{l}
[||\lambda a \to sqrt\ (snd \\
\quad (while\ (\lambda s \to fst\ s < lnArr\ a) \\
\qquad (\lambda s \to \mathbf{let}\ i = fst\ s\ \mathbf{in} \\
\qquad\quad (i + 1, snd\ s + (ixArr\ a\ i \times ixArr\ a\ i))) \\
\qquad (0, 0.0)))||]
\end{array}
$$

from which it is easy to generate C code.

The vector representation makes it easy to define any function where each vector element is computed independently, such as the examples above, vector append ($appVec$) and creating a vector of one element ($uniVec$), but is less well suited to functions with dependencies between elements, such as computing a running sum.

Types and the subformula property help us to guarantee fusion. The subformula property guarantees that all occurrences of Vec must be eliminated, while occurrences of Arr will remain. There are some situations where fusion is not beneficial, notably when an intermediate vector is accessed many times, in which case fusion will cause the elements to be recomputed. An alternative is to materialise the vector as an array with the following function.

$$
\begin{array}{l}
memorise :: Rep\ a \Rightarrow Qt\ (Vec\ a \to Vec\ a) \\
memorise = [||\ \$\$\ toVec \circ save \circ \$\$fromVec||]
\end{array}
$$

Here we interpose $save$, as defined in Section 2.3 to forestall the fusion that would otherwise occur. For example, if

$$
\begin{array}{l}
blur :: Qt\ (Vec\ Float \to Vec\ Float) \\
blur = [||\lambda a \to \$\$zipVec\ (\lambda x\ y \to sqrt\ (x \times y)) \\
\qquad\qquad (\$\$appVec\ (\$\$uniVec\ 0)\ a) \\
\qquad\qquad (\$\$appVec\ a\ (\$\$uniVec\ 0))||]
\end{array}
$$

computes the geometric mean of adjacent elements of a vector, then one may choose to compute either

$$[||\$\$blur \circ \$\$blur||] \quad \text{or} \quad [||\$\$blur \circ \$\$memorise \circ \$\$blur||]$$

with different trade-offs between recomputation and memory use. Strong guarantees for fusion in combination with $memorise$ give the programmer a simple interface which provides powerful optimisations combined with fine control over memory usage.

We have described the application of the subformula to array fusion as based on "pull arrays" (Svenningsson and Axelsson 2012), but the same technique should also apply to other techniques that support array fusion, such as "push arrays" (Claessen et al. 2012).

3. Implementation

The original EDSL Feldspar generates values of a GADT (called Dp in Section 5), with constructs that represent $while$ and manifest arrays similar to those above. A backend then compiles values of type $Dp\ a$ to C code. QDSL Feldspar provides a transformer from $Qt\ a$ to $Dp\ a$, and shares the EDSL Feldspar backend.

The transformer from Qt to Dp performs the following steps.

- In any context where a constant c is not fully applied, it replaces c with $\lambda\overline{x}.\ c\ \ \overline{x}$. It replaces identifiers connected to the type $Maybe$, such as $return$, $(\gg\!=)$, and $maybe$, by their definitions.
- It normalises the term to ensure the subformula property, using the rules of Section 4. The normaliser supports a limited set of types, including tuples, $Maybe$, and Vec.
- It performs simple type inference, which is used to resolve overloading. Overloading is limited to a fixed set of cases, including overloading arithmetic operators at types Int and $Float$.
- It traverses the term, converting Qt to Dp. It checks that only permitted primitives appear in Qt, and translates these to their corresponding representation in Dp. Permitted primitives include: $(==)$, $(<)$, $(+)$, $(\times)$, and similar, plus $while$, $makeArr$, $lenArr$, $ixArr$, and $save$.

An unfortunate feature of typed quasiquotation in GHC is that the implementation discards all type information when creating the representation of a term. Type $Qt\ a$ is a synonym for the type

$$TH.Q\ (TH.TExp\ a)$$

where TH denotes the library for Template Haskell, $TH.Q$ is the quotation monad (used to look up identifiers and generate fresh names), and $TH.TExp\ a$ is the parse tree for a quoted expression returning a value of type a (a wrapper for the type $TH.Exp$ of untyped expressions, with a as a phantom variable). Thus, the translator from $Qt\ a$ to $Dp\ a$ is forced to re-infer all type for subterms, and for this reason we support only limited overloading, and we translate the $Maybe$ monad as a special case rather than supporting overloading for monad operations in general.

The backend performs three transformations over Dp terms before compiling to C. First, common subexpessions are recognised and transformed to **let** bindings. Second, Dp terms are normalised using exactly the same rules used for normalising Qt terms, as described in Section 4. Third, Dp terms are optimised using η contraction for conditionals and arrays:

$$
\begin{array}{rcl}
\mathbf{if}\ L\ \mathbf{then}\ M\ \mathbf{else}\ M & \mapsto & M \\
makeArr\ (lenArr\ M)\ (ixArr\ M) & \mapsto & M
\end{array}
$$

and a restricted form of linear inlining for **let** bindings that preserves the order of evaluation.

Figure 1 lists lines of code, benchmarks used, and performance results. The translator from Dp to C is shared by QDSL and EDSL Feldspar, and listed in a separate column. All five benchmarks run under QDSL and EDSL Felsdpar generate identical C code, up to permutation of independent assignments, with identical compile and run times. The columns for QDSL and EDSL Feldspar give compile and run times for Haskell, while the columns for generated code give compile and run times for the generated C. QDSL compile times are slightly greater than EDSL, and QDSL run times range from identical to four times that of EDSL, the increase being due to normalisation time (our normaliser was not designed to be particularly efficient).

4. The subformula property

This section introduces reduction rules for normalising terms that enforce the subformula property while preserving sharing. The rules adapt to both call-by-need and call-by-value. We work with simple types. The only polymorphism in our examples corresponds to instantiating constants (such as $while$) at different types.

Lines of Haskell code

	shared	unique	total
QDSL Feldspar	3970	1722	5962
EDSL Feldspar	3970	452	4422

Benchmarks

IPGray	Image Processing (Grayscale)
IPBW	Image Processing (Black and White)
FFT	Fast Fourier Transform
CRC	Cyclic Redundancy Check
Window	Average array in a sliding window

Performance

	QDSL Feldspar		EDSL Feldspar		Generated Code	
	Compile	Run	Compile	Run	Compile	Run
IPGray	16.96	0.01	15.06	0.01	0.06	0.39
IPBW	17.08	0.01	14.86	0.01	0.06	0.19
FFT	17.87	0.39	15.79	0.09	0.07	3.02
CRC	17.14	0.01	15.33	0.01	0.05	0.12
Window	17.85	0.02	15.77	0.01	0.06	0.27

Times in seconds; minimum time of ten runs.
Quad-core Intel i7-2640M CPU, 2.80 GHz, 3.7 GiB RAM.
GHC 7.8.3; GCC 4.8.2; Ubuntu 14.04 (64-bit).

Figure 1. Comparison of QDSL and EDSL Feldspar

Types, terms, and values are presented in Figure 2. Let A, B, C range over types, including base types (ι), functions ($A \to B$), products ($A \times B$), and sums ($A+B$). Let L, M, N range over terms, and x, y, z range over variables. Let c range over constants, which are fully applied according to their arity, as discussed below. As constant applications are non-values, we represent literals as free variables. As usual, terms are taken as equivalent up to renaming of bound variables. Write $FV(M)$ for the set of free variables of M, and $N[x := M]$ for capture-avoiding substitution of M for x in N. Let V, W range over values.

Let Γ range over type environments, which pair variables with types, and write $\Gamma \vdash M : A$ to indicate that term M has type A under type environment Γ. Typing rules are standard.

Reduction rules for normalisation are presented in Figure 3. The rules are confluent, so order of application is irrelevant to the final answer, but we break them into three phases to ease the proof of strong normalisation. It is easy to confirm that all of the reduction rules preserve sharing and preserve order of evaluation.

Write $M \mapsto_i N$ to indicate that M reduces to N in phase i. Let F and G range over two different forms of evaluation frame used in Phases 1 and 2 respectively. Write $FV(F)$ for the set of free variables of F, and similarly for G. Reductions are closed under compatible closure.

The normalisation procedure consists of exhaustively applying the reductions of Phase 1 until no more apply, then similarly for Phase 2, and finally for Phase 3. Phase 1 performs let-insertion, naming subterms, along the lines of a translation to A-normal form (Flanagan et al. 1993) or reductions (let.1) and (let.2) in Moggi's metalanguage for monads (Moggi 1991). Phase 2 performs two kinds of reduction: β rules apply when an introduction (construction) is immediately followed by an elimination (deconstruction), and κ rules push eliminators closer to introducers to enable β rules. Phase 3 "garbage collects" unused terms as in the call-by-need lambda calculus (Maraist et al. 1998). Phase 3 should be omitted if the intended semantics of the target language is call-by-value rather than call-by-need.

Every term has a normal form.

PROPOSITION 4.1 (Strong normalisation). *Each of the reduction relations $\mapsto_i$ is confluent and strongly normalising: all $\mapsto_i$ reduction sequences on well-typed terms are finite.*

The only non-trivial proof is for $\mapsto_2$, which can be proved via a standard reducibility argument (see, for example, (Lindley 2007)). If the target language includes general recursion, normalisation should treat the fixpoint operator as an uninterpreted constant.

The *subformulas* of a type are the type itself and its components. For instance, the subformulas of $A \to B$ are itself and the subformulas of A and B. The *proper subformulas* of a type are all its subformulas other than the type itself.

The *subterms* of term are the term itself and its components. For instance, the subterms of $\lambda x.\, N$ are itself and the subterms of N and the subterms of $L\ M$ are itself and the subterms of L and M. The *proper subterms* of a term are all its subterms other than the term itself.

Constants are always fully applied; they are introduced as a separate construct to avoid consideration of irrelevant subformulas and subterms. The type of a constant c of arity k is written

$$c : A_1 \to \cdots \to A_k \to B$$

and its subformulas are itself and A_1, ..., A_k, and B (but not $A_i \to \cdots \to A_k \to B$ for $i > 1$). An application of a constant c of arity k is written

$$c\ M_1\ \cdots\ M_k$$

and its subterms are itself and $M_1, \ldots, M_k$ (but not $c\ M_1\ \cdots\ M_j$ for $j < k$). Free variables are equivalent to constants of arity zero.

Terms in normal form satisfy the subformula property.

PROPOSITION 4.2 (Subformula property). *If $\Gamma \vdash M : A$ and M is in normal form, then every subterm of M has a type that is either a subformula of A, a subformula of a type in Γ, or a subformula of the type of a constant in M.*

The proof follows the lines of Prawitz (1965). The differences are that we have introduced fully applied constants (to enable the sharpened subformula property, below), and that our reduction rules introduce **let**, in order to ensure sharing is preserved.

Normalisation may lead to an exponential or worse blow up in the size of a term, for instance when there are nested **case** expressions. The benchmarks in Section 3 do not suffer from blow up, but it may be a problem in some contexts. Normalisation may be controlled by introduction of uninterpreted constants, as in Section 2.3. Further work is needed to understand when complete normalisation is desirable and when it is problematic.

Examination of the proof in Prawitz (1965) shows that in fact normalisation achieves a sharper property.

PROPOSITION 4.3 (Sharpened subformula). *If $\Gamma \vdash M : A$ and M is in normal form, then every proper subterm of M that is not a free variable or a subterm of a constant application has a type that is a proper subformula of A or a proper subformula of a type in Γ.*

We believe we are the first to formulate the sharpened version.

The sharpened subformula property says nothing about the types of subterms of constant applications, but such types are immediately apparent by recursive application of the sharpened subformula property. Given a subterm that is a constant application $c\ \overline{M}$, where c has type $\overline{A} \to B$, then the subterm itself has type B, each subterm M_i has type A_i, and every proper subterm of M_i that is not a free variable of M_i or a subterm of a constant application has a type that is a proper subformula of A_i or a proper subformula of the type of one of its free variables.

In Section 2, we require that every top-level term passed to *qdsl* is suitable for translation to C after normalisation, and any DSL translating to a *first-order* language must impose a similar requirement. One might at first guess the required property is that every

$$
\begin{array}{llll}
\text{Types} & A, B, C & ::= & \iota \mid A \to B \mid A \times B \mid A + B \\
\text{Terms} & L, M, N & ::= & x \mid c\ \overline{M} \mid \lambda x.\, N \mid L\ M \mid \mathbf{let}\ x = M\ \mathbf{in}\ N \mid (M, N) \mid \mathbf{fst}\ L \mid \mathbf{snd}\ L \mid \\
 & & & \mathbf{inl}\ M \mid \mathbf{inr}\ N \mid \mathbf{case}\ L\ \mathbf{of}\ \{\mathbf{inl}\ x.\, M;\ \mathbf{inr}\ y.\, N\} \\
\text{Values} & V, W & ::= & x \mid \lambda x.\, N \mid (V, W) \mid \mathbf{inl}\ V \mid \mathbf{inr}\ W
\end{array}
$$

Figure 2. Types and Terms

Phase 1 (let-insertion)

$$F ::= [\,]\ M \mid V\ [\,] \mid ([\,], M) \mid (V, [\,]) \mid \mathbf{fst}\ [\,] \mid \mathbf{snd}\ [\,] \mid \mathbf{inl}\ [\,] \mid \mathbf{inr}\ [\,] \mid \mathbf{case}\ [\,]\ \mathbf{of}\ \{\mathbf{inl}\ x.\, M;\ \mathbf{inr}\ y.\, N\}$$

$$(let) \quad F[M] \quad \mapsto_1 \quad \mathbf{let}\ x = M\ \mathbf{in}\ F[x], \quad x \text{ fresh, } M \text{ not a value}$$

Phase 2 (symbolic evaluation)

$$G ::= \mathbf{let}\ x = [\,]\ \mathbf{in}\ N$$

$$
\begin{array}{llll}
(\kappa.let) & G[\mathbf{let}\ x = M\ \mathbf{in}\ N] & \mapsto_2 & \mathbf{let}\ x = M\ \mathbf{in}\ G[N], \qquad x \notin FV(G) \\
(\kappa.case) & G[\mathbf{case}\ V\ \mathbf{of}\ \{\mathbf{inl}\ x.\, M;\ \mathbf{inr}\ y.\, N\}] & \mapsto_2 & \mathbf{case}\ V\ \mathbf{of}\ \{\mathbf{inl}\ x.\, G[M];\ \mathbf{inr}\ y.\, G[N]\}, \quad x, y \notin FV(G) \\
(\beta.\to) & (\lambda x.\, N)\ V & \mapsto_2 & N[x := V] \\
(\beta.\times_1) & \mathbf{fst}\ (V, W) & \mapsto_2 & V \\
(\beta.\times_2) & \mathbf{snd}\ (V, W) & \mapsto_2 & W \\
(\beta.+_1) & \mathbf{case}\ (\mathbf{inl}\ V)\ \mathbf{of}\ \{\mathbf{inl}\ x.\, M;\ \mathbf{inr}\ y.\, N\} & \mapsto_2 & M[x := V] \\
(\beta.+_2) & \mathbf{case}\ (\mathbf{inr}\ W)\ \mathbf{of}\ \{\mathbf{inl}\ x.\, M;\ \mathbf{inr}\ y.\, N\} & \mapsto_2 & N[y := W] \\
(\beta.let) & \mathbf{let}\ x = V\ \mathbf{in}\ N & \mapsto_2 & N[x := V]
\end{array}
$$

Phase 3 (garbage collection)

$$(need) \quad \mathbf{let}\ x = M\ \mathbf{in}\ N \quad \mapsto_3 \quad N, \quad x \notin FV(N)$$

Figure 3. Normalisation Rules

subterm is *representable*, in the sense introduced in Section 2.1, but this is not quite right. The top-level term is a function from a representable type to a representable type, and the constant *while* expects subterms of type $s \to \mathit{Bool}$ and $s \to s$, where the state s is representable. Fortunately, the property required is not hard to formulate in a general way, and is easy to ensure by applying the sharpened subformula property.

Take the representable types to be any set closed under subformulas that does not include function types. We introduce a variant of the usual notion of *rank* of a type, with respect to a notion of representability. A term of type $A \to B$ has rank $\min(m + 1, n)$ where m is the rank of A and n is the rank of B, while a term of representable type has rank 0. We say a term is *first-order* when every subterm is either representable, or is of the form $\lambda\overline{x}.\, N$ where each bound variable and the body is of representable type.

The following characterises translation to a first-order language.

PROPOSITION 4.4 (First-order). *Consider a term of rank* 1*, where every free variable has rank* 0 *and every constant has rank at most* 2*. Then the term normalises to a term that is first-order.*

The property follows immediately by observing that any term L of rank 1 can be rewritten to the form $\lambda\overline{y}.\, (L\ \overline{y})$ where each bound variable and the body has representable type, and then normalising and applying the sharpened subformula property.

In QDSL Feldspar, *while* is a constant with type of rank 2 and other constants have types of rank 1. Section 2.6 gives an example of a normalised term. By the proposition, each subterm has a representable type (boolean, integer, float, or a pair of an integer and float) or is a lambda abstraction with bound variables and body of representable type; and it is this property which ensures it is easy to generate C code from the term.

5. Feldspar as an EDSL

This section reviews the combination of deep and shallow embeddings required to implement Feldspar as an EDSL, and considers the trade-offs between the QDSL and EDSL approaches. Much of this section reprises Svenningsson and Axelsson (2012).

The top-level function of EDSL Feldspar has the type:

$$\mathit{edsl} :: (\mathit{Rep}\ a, \mathit{Rep}\ b) \Rightarrow (\mathit{Dp}\ a \to \mathit{Dp}\ b) \to C$$

Here $\mathit{Dp}\ a$ is the deep representation of a term of type a. The deep representation is described in detail in Section 5.3 below, and is chosen to be easy to translate to C. As before, type C represents code in C, and type class *Rep* restricts to representable types.

5.1 A first example

Here is the power function of Section 2.2, written as an EDSL:

```
power :: Int → Dp Float → Dp Float
power n x =
  if n < 0 then
    x .==. 0 ? (0, 1 / power (−n) x)
  else if n == 0 then
    1
  else if even n then
    let y = power (n div 2) x in y × y
  else
    x × power (n − 1) x
```

Type $Q\ (\mathit{Float} \to \mathit{Float})$ in the QDSL variant becomes the type $\mathit{Dp}\ \mathit{Float} \to \mathit{Dp}\ \mathit{Float}$ in the EDSL variant, meaning that $\mathit{power}\ n$ accepts a representation of the argument and returns a representation of that argument raised to the n'th power.

In the EDSL variant, no quotation is required, and the code looks almost—but not quite!—like an unstaged version of power,

but with different types. Clever encoding tricks, explained later, permit declarations, function calls, arithmetic operations, and numbers to appear the same whether they are to be executed at generation-time or run-time. However, as explained later, comparison and conditionals appear differently depending on whether they are to be executed at generation-time or run-time, using $M == N$ and **if** L **then** M **else** N for the former but $M \mathbin{.==.} N$ and $L \mathbin{?} (M, N)$ for the latter.

Invoking $edsl\ (power\ (-6))$ generates code to raise a number to its -6 power. Evaluating $power\ (-6)\ u$, where u is a term representing a variable of type $Dp\ Float$, yields the following:

$$\begin{array}{l}(u \mathbin{.==.} 0) \mathbin{?} (0,\\ \quad 1 \mathbin{/} ((u \times ((u \times 1) \times (u \times 1))) \times\\ \qquad (u \times ((u \times 1) \times (u \times 1)))))\end{array}$$

Applying common-subexpression elimination permits recovering the sharing structure.

$$\begin{array}{r|l} v & (u \times 1) \\ w & u \times (v \times v) \\ \text{top} & (u \mathbin{.==.} 0) \mathbin{?} (0, 1 \mathbin{/} (w \times w)) \end{array}$$

From the above, it is easy to generate the final C code, which is identical to that in Section 2.2.

Here are points of comparison between the two approaches.

- A function $a \to b$ is embedded in QDSL as $Qt\ (a \to b)$, a representation of a function, and in EDSL as $Dp\ a \to Dp\ b$, a function between representations.
- QDSL enables the host and embedded languages to appear identical. In contrast, in Haskell, EDSL requires some term forms, such as comparison and conditionals, to differ between the host and embedded languages. Other languages, notably Scala Virtualised (Rompf et al. 2013), may support more general overloading that allows even comparison and conditionals to be identical.
- QDSL requires syntax to separate quoted and unquoted terms. In contrast, EDSL permits the host and embedded languages to intermingle seamlessly. Depending on your point of view, explicit quotation syntax may be considered as an unnecessary distraction or as drawing a useful distinction between generation-time and run-time. If one takes the former view, the type-based approach to quotation found in C# and Scala might be preferred.
- QDSL may share the same representation for quoted terms across a range of applications; the quoted language is the host language, and does not vary with the specific domain. In contrast, EDSL typically develops custom shallow and deep embeddings for each application; a notable exception is the LMS and Delite frameworks for Scala, which provide a deep embedding shared across several disparate DSLs (Sujeeth et al. 2013).
- QDSL yields an unwieldy term that requires normalisation. In contrast, EDSL yields the term in normalised form in this case, though there are other situations where a normaliser is required (see Section 5.2).
- QDSL requires traversing the quoted term to ensure it only mentions permitted identifiers. In contrast, EDSL guarantees that if a term has the right type it will translate to the target. If the requirement to eyeball code to ensure only permitted identifiers are used is considered too onerous, it should be easy to build a preprocessor that checks this property. For example, in Haskell, it is possible to incorporate such a preprocessor using MetaHaskell (Mainland 2012).
- Since QDSLs may share the same quoted terms across a range of applications, the cost of building a normaliser or a preprocessor might be amortised across multiple QDSLs for a single language. In the conclusion, we consider the design of a tool for building QDSLs that uses a shared normaliser and preprocessor.
- Once the deep embedding or the normalised quoted term is produced, generating the domain-specific code is similar for both approaches.

5.2 A second example

In Section 2.4, we exploited the $Maybe$ type to refactor the code.

In EDSL, we must use a new type, where $Maybe$, $Nothing$, $Just$, and $maybe$ become Opt, $none$, $some$, and $option$, and $return$ and $(\gg\!=)$ are similar to before.

$$\begin{array}{l}\textbf{type}\ Opt\ a\\ none\ \ :: Undef\ a \Rightarrow Opt\ a\\ some\ \ :: a \to Opt\ a\\ return :: a \to Opt\ a\\ (\gg\!=)\ \ :: Opt\ a \to (a \to Opt\ b) \to Opt\ b\\ option :: (Undef\ a, Undef\ b) \Rightarrow\\ \qquad\qquad b \to (a \to b) \to Opt\ a \to b\end{array}$$

Type class $Undef$ is explained in Section 5.6, and details of type Opt are given in Section 5.7.

Here is the refactored code.

$$\begin{array}{l}power' :: Int \to Dp\ Float \to Opt\ (Dp\ Float)\\ power'\ n\ x =\\ \quad \textbf{if}\ n < 0\ \textbf{then}\\ \quad\quad (x \mathbin{.==.} 0) \mathbin{?} (none,\\ \quad\quad\quad\quad \textbf{do}\ y \leftarrow power'\ (-n)\ x\\ \quad\quad\quad\quad\quad return\ (1 \mathbin{/} y))\\ \quad \textbf{else if}\ n == 0\ \textbf{then}\\ \quad\quad return\ 1\\ \quad \textbf{else if}\ even\ n\ \textbf{then}\\ \quad\quad \textbf{do}\ y \leftarrow power'\ (n\ \text{div}\ 2)\ x\\ \quad\quad\quad return\ (y \times y)\\ \quad \textbf{else}\\ \quad\quad \textbf{do}\ y \leftarrow power'\ (n-1)\ x\\ \quad\quad\quad return\ (x \times y)\\ power'' :: Int \to Dp\ Float \to Dp\ Float\\ power''\ n\ x = option\ 0\ (\lambda y \to y)\ (power'\ n\ x)\end{array}$$

The term of type $Dp\ Float$ generated by evaluating $power\ (-6)\ x$ is large and unscrutable:

$$\begin{array}{l}(((fst\ ((x == 0.0) \mathbin{?} (((False \mathbin{?} (True, False)), (False \mathbin{?}\\ (undef, undef))), (True, (1.0 \mathbin{/} ((x \times ((x \times 1.0) \times (x \times\\ 1.0))) \times (x \times ((x \times 1.0) \times (x \times 1.0)))))))))) \mathbin{?} (True,\\ False)) \mathbin{?} (((fst\ ((x == 0.0) \mathbin{?} (((False \mathbin{?} (True, False)),\\ (False \mathbin{?} (undef, undef))), (True, (1.0 \mathbin{/} ((x \times ((x \times 1.0) \times\\ (x \times 1.0))) \times (x \times ((x \times 1.0) \times (x \times 1.0)))))))))) \mathbin{?} ((snd\\ ((x == 0.0) \mathbin{?} (((False \mathbin{?} (True, False)), (False \mathbin{?} (undef,\\ undef))), (True, (1.0 \mathbin{/} ((x \times ((x \times 1.0) \times (x \times 1.0))) \times\\ (x \times ((x \times 1.0) \times (x \times 1.0)))))))))), undef)), 0.0))\end{array}$$

Before, evaluating $power$ yielded a term essentially in normal form. However, here rewrite rules need to be repeatedly applied, as described in Section 3. After applying these rules, common subexpression elimination yields the same structure as in the previous subsection, from which the same C code is generated.

Here we have described normalisation via rewriting, but some EDSLs achieve normalisation via smart constructors, which ensure deep terms are always in normal form (Rompf 2012); the two techniques are roughly equivalent.

Hence, an advantage of the EDSL approach—that it generates terms essentially in normal form—turns out to apply sometimes but not others. It appears to often work for functions and products, but to fail for sums. In such situations, separate normalisation is required. This is one reason why we do not consider normalisation as required by QDSL to be particularly onerous.

Here are points of comparison between the two approaches.

- Both QDSL and EDSL can exploit notational conveniences in the host language. The example here exploits Haskell **do** notation; the embedding of SQL in F# by Cheney et al. (2013) expoits F# sequence notation. For EDSL, exploiting **do** notation just requires instantiating *return* and $(\gg\!=)$ correctly. For QDSL, it is also necessary for the translator to recognise and expand **do** notation and to substitute appropriate instances of *return* and $(\gg\!=)$.
- As this example shows, sometimes both QDSL and EDSL may require normalisation. As mentioned previously, for QDSLs the cost of building a normaliser might be amortised across several applications. In constrast, each EDSL usually has a distinct deep representation and so requires a distinct normaliser.

5.3 The deep embedding

Recall that a value of type $Dp\ a$ represents a term of type a, and is called a deep embedding.

```
data Dp a where
  LitB      :: Bool → Dp Bool
  LitI      :: Int → Dp Int
  LitF      :: Float → Dp Float
  If        :: Dp Bool → Dp a → Dp a → Dp a
  While     :: (Dp a → Dp Bool) →
               (Dp a → Dp a) → Dp a → Dp a
  Pair      :: Dp a → Dp b → Dp (a, b)
  Fst       :: Rep b ⇒ Dp (a, b) → Dp a
  Snd       :: Rep a ⇒ Dp (a, b) → Dp b
  Prim1     :: Rep a ⇒ String → Dp a → Dp b
  Prim2     :: (Rep a, Rep b) ⇒
               String → Dp a → Dp b → Dp c
  MkArr     :: Dp Int → (Dp Int → Dp a) → Dp (Arr a)
  LnArr     :: Rep a ⇒ Dp (Arr a) → Dp Int
  IxArr     :: Dp (Arr a) → Dp Int → Dp a
  Save      :: Dp a → Dp a
  Let       :: Rep a ⇒ Dp a → (Dp a → Dp b) → Dp b
  Variable  :: String → Dp a
```

Type Dp represents a low level, pure functional language with a straightforward translation to C. It uses higher-order abstract syntax (HOAS) to represent constructs with variable binding Pfenning and Elliot (1988). Our code obeys the invariant that we only write $Dp\ a$ when $Rep\ a$ holds, that is, when type a is representable.

The deep embedding has boolean, integer, and floating point literals, conditionals, while loops, pairs, primitives, arrays, and special-purpose constructs to disable normalisation, for let binding, and for variables. Constructs $LitB$, $LitI$, $LitF$ build literals; If builds a conditional. $While$ corresponds to $while$ in Section 2.5; $Pair$, Fst, and Snd build and decompose pairs; $Prim1$ and $Prim2$ represent primitive operations, where the string is the name of the operation; $MkArr$, $LnArr$, and $IxArr$ correspond to the array operations in Section 2.6; $Save$ corresponds to $save$ in Section 2.3; Let corresponds to let binding, and $Variable$ is used when translating HOAS to C code.

5.4 Class Syn

We introduce a type class Syn that allows us to convert shallow embeddings to and from deep embeddings.

```
class Rep (Internal a) ⇒ Syn a where
  type Internal a
  toDp   :: a → Dp (Internal a)
  fromDp :: Dp (Internal a) → a
```

Type $Internal$ is a GHC type family (Chakravarty et al. 2005). Functions $toDp$ and $fromDp$ translate between the shallow embedding a and the deep embedding $Dp\ (Internal\ a)$.

The first instance of Syn is Dp itself, and is straightforward.

```
instance Rep a ⇒ Syn (Dp a) where
  type Internal (Dp a) = a
  toDp   = id
  fromDp = id
```

Our representation of a run-time $Bool$ will have type $Dp\ Bool$ in both the deep and shallow embeddings, and similarly for Int and $Float$.

We do not code the target language using its constructs directly. Instead, for each constructor we define a corresponding "smart constructor" using class Syn.

```
true, false :: Dp Bool
true = LitB True
false = LitB False

(?) :: Syn a ⇒ Dp Bool → (a, a) → a
c ? (t, e) = fromDp (If c (toDp t) (toDp e))

while :: Syn a ⇒ (a → Dp Bool) → (a → a) → a → a
while c b i = fromDp (While (c ∘ fromDp)
                            (toDp ∘ b ∘ fromDp)
                            (toDp i))
```

Numbers are made convenient to manipulate via overloading.

```
instance Num (Dp Int) where
  a + b = Prim2 "(+)" a b
  a − b = Prim2 "(-)" a b
  a × b = Prim2 "(*)" a b
  fromInteger a = LitI (fromInteger a)
```

With this declaration, $1 + 2 :: Dp\ Int$ evaluates to

$$Prim2\ \texttt{"(+)"}\ (LitI\ 1)\ (LitI\ 2),$$

permitting code executed at generation-time and run-time to appear identical. A similar declaration works for $Float$.

Comparison also benefits from smart constructors.

```
(.==.) :: (Syn a, Eq (Internal a)) ⇒ a → a → Dp Bool
a .==. b = Prim2 "(==)" (toDp a) (toDp b)

(.<.) :: (Syn a, Ord (Internal a)) ⇒ a → a → Dp Bool
a .<. b = Prim2 "(<)" (toDp a) (toDp b)
```

Overloading cannot apply here, because Haskell requires $(==)$ return a result of type $Bool$, while $(.==.)$ returns a result of type $Dp\ Bool$, and similarly for $(.<.)$.

Here is how to compute the minimum of two values.

```
minim      :: (Syn a, Ord (Internal a)) ⇒ a → a → a
minim x y = (x .<. y) ? (x, y)
```

5.5 Embedding pairs

Host language pairs in the shallow embedding correspond to target language pairs in the deep embedding.

```
instance (Syn a, Syn b) ⇒ Syn (a, b) where
  type Internal (a, b) = (Internal a, Internal b)
  toDp (a, b) = Pair (toDp a) (toDp b)
  fromDp p    = (fromDp (Fst p), fromDp (Snd p))
```

This permits us to manipulate pairs as normal, with (a, b), *fst* a, and *snd* a. Argument p is duplicated in the definition of *fromDp*, which may require common subexpression elimination as discussed in Section 5.1.

We have now developed sufficient machinery to define a *for* loop in terms of a *while* loop.

```
for :: Syn a ⇒ Dp Int → a → (Dp Int → a → a) → a
for n s₀ b = snd (while (λ(i, s) → i .<. n)
                        (λ(i, s) → (i + 1, b i s))
                        (0, s₀))
```

The state of the *while* loop is a pair consisting of a counter and the state of the *for* loop. The body b of the *for* loop is a function that expects both the counter and the state of the *for* loop. The counter is discarded when the loop is complete, and the final state of the *for* loop returned.

Thanks to our machinery, the above definition uses only ordinary Haskell pairs. The condition and body of the *while* loop pattern match on the state using ordinary pair syntax, and the initial state is constructed as an ordinary pair.

5.6 Embedding undefined

For the next section, which defines an analogue of the *Maybe* type, it will prove convenient to work with types which have a distinguished value at each type, which we call *undef*.

It is straightforward to define a type class *Undef*, where type a belongs to *Undef* if it belongs to *Syn* and it has an undefined value.

```
class Syn a ⇒ Undef a where
  undef :: a

instance Undef (Dp Bool) where
  undef = false

instance Undef (Dp Int) where
  undef = 0

instance Undef (Dp Float) where
  undef = 0

instance (Undef a, Undef b) ⇒ Undef (a, b) where
  undef = (undef, undef)
```

For example,

```
(/#)   :: Dp Float → Dp Float → Dp Float
x /# y = (y .==. 0) ? (undef, x / y)
```

behaves as division, save that when the divisor is zero it returns the undefined value of type *Float*, which is also zero.

Svenningsson and Axelsson (2012) claim that it is not possible to support *undef* without changing the deep embedding, but here we have defined *undef* entirely as a shallow embedding. (It appears they underestimated the power of their own technique!)

5.7 Embedding option

We now explain in detail the *Opt* type seen in Section 2.4.

The deep-and-shallow technique represents deep embeddding *Dp* (a, b) by shallow embedding $(Dp\ a, Dp\ b)$. Hence, it is tempting to represent *Dp* (*Maybe* a) by *Maybe* (*Dp* a), but this cannot work, because *fromDp* would have to decide at generation-time whether to return *Just* or *Nothing*, but which to use is not known until run-time.

Instead, Svenningsson and Axelsson (2012) represent values of type *Maybe* a by the type *Opt′* a, which pairs a boolean with a value of type a. For a value corresponding to *Just* x, the boolean is true and the value is x, while for one corresponding to *Nothing*, the boolean is false and the value is *undef*. We define *some′*, *none′*, and *option′* as the analogues of *Just*, *Nothing*, and *maybe*. The *Syn* instance is straightforward, mapping options to and from the pairs already defined for *Dp*.

```
data Opt′ a = Opt′ {def :: Dp Bool, val :: a}

instance Syn a ⇒ Syn (Opt′ a) where
  type Internal (Opt′ a) = (Bool, Internal a)
  toDp (Opt′ b x) = Pair b (toDp x)
  fromDp p        = Opt′ (Fst p) (fromDp (Snd p))

some′       ::  a → Opt′ a
some′ x     =   Opt′ true x

none′       ::  Undef a ⇒ Opt′ a
none′       =   Opt′ false undef

option′     ::  Syn b ⇒ b → (a → b) → Opt′ a → b
option′ d f o = def o ? (f (val o), d)
```

The next obvious step is to define a suitable monad over the type *Opt′*. The natural definitions to use are as follows:

```
return   :: a → Opt′ a
return x = some′ x

(≫=)     :: (Undef b) ⇒ Opt′ a → (a → Opt′ b) → Opt′ b
o ≫= g   = Opt′ (def o ? (def (g (val o)), false))
                (def o ? (val (g (val o)), undef))
```

However, this adds type constraint *Undef* b to the type of $(\gg\!=)$, which is not permitted. The need to add such constraints often arises, and has been dubbed the constrained-monad problem (Hughes 1999; Svenningsson and Svensson 2013). We solve it with a trick due to Persson et al. (2011).

We introduce a continuation-passing style (CPS) type, *Opt*, defined in terms of *Opt′*. It is straightforward to define *Monad* and *Syn* instances, operations to lift the representation type to lift and lower one type to the other, and to lift *some*, *none*, and *option* to the CPS type. The *lift* operation is closely related to the $(\gg\!=)$ operation we could not define above; it is properly typed, thanks to the type constraint on b in the definition of *Opt* a.

```
newtype Opt a =
  O {unO :: ∀b. Undef b ⇒ ((a → Opt′ b) → Opt′ b)}

instance Monad Opt where
  return x = O (λg → g x)
  m ≫= k   = O (λg → unO m (λx → unO (k x) g))

instance Undef a ⇒ Syn (Opt a) where
  type Internal (Opt a) = (Bool, Internal a)
  fromDp = lift ∘ fromDp
  toDp   = toDp ∘ lower

lift :: Opt′ a → Opt a
lift o =  O (λg → Opt′ (def o ? (def (g (val o)), false))
                       (def o ? (val (g (val o)), undef)))

lower :: Undef a ⇒ Opt a → Opt′ a
lower m = unO m some′

none :: Undef a ⇒ Opt a
none = lift none′

some :: a → Opt a
some a = lift (some′ a)
```

$$option :: (Undef\ a, Syn\ b) \Rightarrow b \to (a \to b) \to Opt\ a \to b$$
$$option\ d\ f\ o = option'\ d\ f\ (lower\ o)$$

These definitions support the EDSL code presented in Section 5.2.

5.8 Embedding vector

Recall that values of type $Array$ are created by construct $MkArr$, while $LnArr$ extracts the length and $IxArr$ fetches the element at the given index. Corresponding to the deep embedding $Array$ is a shallow embedding Vec.

$$\textbf{data}\ Vec\ a = Vec\ (Dp\ Int)\ (Dp\ Int \to a)$$
$$\textbf{instance}\ Syn\ a \Rightarrow Syn\ (Vec\ a)\ \textbf{where}$$
$$\quad \textbf{type}\ Internal\ (Vec\ a) = Array\ Int\ (Internal\ a)$$
$$\quad toDp\ (Vec\ n\ g) \quad = MkArr\ n\ (toDp \circ g)$$
$$\quad fromDp\ a \qquad\quad = Vec\ (LnArr\ a)\ (fromDp \circ IxArr\ a)$$
$$\textbf{instance}\ Functor\ Vec\ \textbf{where}$$
$$\quad fmap\ f\ (Vec\ n\ g) = Vec\ n\ (f \circ g)$$

Constructor Vec resembles Arr, but the former constructs a high-level representation of the array and the latter an actual array. It is straightforward to make Vec an instance of $Functor$.

It is easy to define operations on vectors, including combining corresponding elements of two vectors, summing the elements of a vector, dot product of two vectors, and norm of a vector.

$$zipVec\ \ :: (Syn\ a, Syn\ b) \Rightarrow$$
$$\qquad (a \to b \to c) \to Vec\ a \to Vec\ b \to Vec\ c$$
$$zipVec\ f\ (Vec\ m\ g)\ (Vec\ n\ h)$$
$$\qquad = Vec\ (m\ `minim`\ n)\ (\lambda i \to f\ (g\ i)\ (h\ i))$$
$$sumVec :: (Syn\ a, Num\ a) \Rightarrow Vec\ a \to a$$
$$sumVec\ (Vec\ n\ g)$$
$$\qquad = for\ n\ 0\ (\lambda i\ x \to x + g\ i)$$
$$dotVec\ \ :: (Syn\ a, Num\ a) \Rightarrow Vec\ a \to Vec\ a \to a$$
$$dotVec\ u\ v = sumVec\ (zipVec\ (\times)\ u\ v)$$
$$normVec\ \ :: Vec\ (Dp\ Float) \to Dp\ Float$$
$$normVec\ v = sqrt\ (dotVec\ v\ v)$$

Invoking $edsl$ on

$$normVec \circ toVec$$

generates C code to normalise a vector. If we used a top-level function of type $(Syn\ a, Syn\ b) \Rightarrow (a \to b) \to C$, then it would insert the $toVec$ coercion automatically.

This style of definition again provides fusion. For instance:

$$\begin{array}{ll} & dotVec\ (Vec\ m\ g)\ (Vec\ n\ h) \\ = & sumVec\ (zipVec\ (\times)\ (Vec\ m\ g)\ (Vec\ n\ h) \\ = & sumVec\ (Vec\ (m\ `minim`\ n)\ (\lambda i \to g\ i \times h\ i) \\ = & for\ (m\ `minim`\ n)\ (\lambda i\ x \to x + g\ i \times h\ i) \end{array}$$

Indeed, we can see that by construction that whenever we combine two primitives the intermediate vector is always eliminated.

The type class Syn enables conversion between types Arr and Vec. Hence for EDSL, unlike QDSL, explicit calls $toVec$ and $fromVec$ are not required. Invoking $edsl\ normVec$ produces the same C code as in Section 2.6.

As with QDSL, there are some situations where fusion is not beneficial. We may materialise a vector as an array with the following function.

$$memorise :: Syn\ a \Rightarrow Vec\ a \to Vec\ a$$
$$memorise = fromDp \circ Save \circ toDp$$

Here we interpose $Save$ to forestall the fusion that would otherwise occur. For example, if

$$blur :: Syn\ a \Rightarrow Vec\ a \to Vec\ a$$
$$blur\ v = zipVec\ (\lambda x\ y \to sqrt\ (x \times y))$$
$$\qquad (appVec\ a\ (uniVec\ 0))$$
$$\qquad (appVec\ (uniVec\ 0)\ a)$$

computes the geometric mean of adjacent elements of a vector, then one may choose to compute either

$$blur \circ blur \quad \text{or} \quad blur \circ memorise \circ blur$$

with different trade-offs between recomputation and memory use.

QDSL forces all conversions to be written out, while EDSL silently converts between representations; following the pattern that QDSL is more explicit, while EDSL is more compact. For QDSL it is the subformula property which guarantees that all intermediate uses of Vec are eliminated, while for EDSL this is established by operational reasoning on the behaviour of the type Vec.

6. Related work

DSLs have a long and rich history (Bentley 1986). An early use of quotation in programming is Lisp (McCarthy 1960), and perhaps the first application of quotation to domain-specific languages is Lisp macros (Hart 1963).

This paper uses Haskell, which has been widely used for EDSLs (Hudak 1997). We constrast QDSL with an EDSL technique that combines deep and shallow embedding, as described by Svenningsson and Axelsson (2012), and as used in several Haskell EDSLs including Feldspar (Axelsson et al. 2010), Obsidian (Svensson et al. 2011), Nikola (Mainland and Morrisett 2010), Hydra (Giorgidze and Nilsson 2011), and Meta-Repa (Ankner and Svenningsson 2013).

O'Donnell (1993) identified loss of sharing in the context of embedded circuit descriptions. Claessen and Sands (1999) extended Haskell to support observable sharing. Gill (2009) proposes library features that support sharing without need to extend the language.

A proposition-as-types principle for quotation as a modal logic was proposed by Davies and Pfenning (2001). As they note, their technique has close connections to two-level languages (Nielson and Nielson 2005) and partial evaluation (Jones et al. 1993).

Other approaches to DSL that make use of quotation include C# and F# versions of LINQ (Meijer et al. 2006; Syme 2006) and Scala Lightweight Modular Staging (LMS) (Rompf and Odersky 2010). Scala LMS exploits techniques found in both QDSL (quotation and normalisation) and EDSL (combining shallow and deep embeddings), see Rompf et al. (2013), and exploits reuse to allow multiple DSLs to share infrastructure see Sujeeth et al. (2013).

The underlying idea for QDSLs was established for F# LINQ by Cheney et al. (2013), and extended to nested results by Cheney et al. (2014b). Related work combines language-integrated query with effect types (Cooper 2009; Lindley and Cheney 2012). Cheney et al. (2014a) compare approaches based on quotation and effects.

7. Conclusion

> A good idea can be much better than a new one.
>
> – Gerard Berry

We have compared QDSLs and EDSLs, arguing that QDSLs offer competing expressiveness and efficiency.

The subformula property may have applications in DSLs other that QDSLs. For instance, after Section 5.7 of this paper was drafted, it occurred to us that a different approach would be to extend type Dp with constructs for type $Maybe$. So long as type $Maybe$ does not appear in the input or output of the program, a normaliser that ensures the subformula property could guarantee that C code for such constructs need never be generated.

Rather than building a special-purpose tool for each QDSL, it should be possible to design a single tool for each host language. Our next step is to design a QDSL library for Haskell that restores

the type information for quasi-quotations currencly discarded by GHC and uses this to support type classes and overloading in full, and to supply a more general normaliser. Such a tool would subsume the special-purpose translator from Qt to Dp described at the beginning of Section 3, and lift most of its restrictions.

> These forty years now I've been speaking in prose without knowing it! — Moliere

Like Molière's Monsieur Jourdain, many of us have used QDSLs for years, if not by that name. DSL via quotation lies at the heart of Lisp macros, Microsoft LINQ, and Scala LMS, to name but three. By naming the concept and drawing attention to the benefits of normalisation and the subformula propety, we hope to help the concept to prosper for years to come.

Acknowledgement Najd is supported by a Google Europe Fellowship in Programming Technology. Svenningsson is a SICSA Visiting Fellow and is funded by a HiPEAC collaboration grant, and by the Swedish Foundation for Strategic Research under grant RawFP. Lindley and Wadler are funded by EPSRC Grant EP/K034413/1.

References

J. Ankner and J. Svenningsson. An EDSL approach to high performance Haskell programming. In *Haskell*, 2013.

E. Axelsson et al. Feldspar: A domain specific language for digital signal processing algorithms. In *MEMOCODE*, 2010.

J. Bentley. Programming pearls: Little languages. *CACM*, 29(8), 1986.

M. Chakravarty, G. Keller, S. Peyton Jones, and S. Marlow. Associated types with class. In *POPL*. ACM, 2005.

J. Cheney, S. Lindley, and P. Wadler. A practical theory of language-integrated query. In *ICFP*. ACM, 2013.

James Cheney, Sam Lindley, Gabriel Radanne, and Philip Wadler. Effective quotation. In *PEPM*, 2014a.

James Cheney, Sam Lindley, and Philip Wadler. Query shredding. In *SIGMOD*, 2014b.

K. Claessen and D. Sands. Observable sharing for functional circuit description. In *ASIAN*. Springer, 1999.

K. Claessen, M. Sheeran, and B. Svensson. Expressive array constructs in an embedded GPU kernel programming language. In *DAMP*. ACM, 2012.

Ezra Cooper. The script-writer's dream. In *DBPL*, 2009.

Rowan Davies and Frank Pfenning. A modal analysis of staged computation. *JACM*, 48(3), 2001.

Cormac Flanagan, Amr Sabry, Bruce F Duba, and Matthias Felleisen. The essence of compiling with continuations. In *PLDI*. ACM, 1993.

Gerhard Gentzen. Untersuchungen über das logische schließen. i. *Mathematische zeitschrift*, 39(1), 1935.

A. Gill. Type-safe observable sharing in Haskell. In *Haskell*, 2009.

G. Giorgidze and H. Nilsson. Embedding a functional hybrid modelling language in Haskell. In *IFL*. Springer, 2011.

Timothy P. Hart. MACRO definitions for LISP. Technical Report AIM-057, MIT, 1963.

W. A. Howard. The formulae-as-types notion of construction. In *To H. B. Curry*, pages 479–491. Academic Press, 1980.

P. Hudak. Domain-specific languages. *Handbook of Programming Languages*, 3, 1997.

J. Hughes. Restricted data types in Haskell. In *Haskell*, 1999.

Neil D Jones, Carsten K Gomard, and Peter Sestoft. *Partial evaluation and automatic program generation*. Prentice Hall, 1993.

S. Lindley. Extensional rewriting with sums. In *TLCA*. Springer, 2007.

Sam Lindley and James Cheney. Row-based effect types for database integration. In Benjamin C. Pierce, editor, *TLDI*. ACM, 2012.

G. Mainland. Explicitly heterogeneous metaprogramming with MetaHaskell. In *ICFP*. ACM, 2012.

G. Mainland and G. Morrisett. Nikola: embedding compiled GPU functions in Haskell. In *Haskell*. ACM, 2010.

J. Maraist, M. Odersky, and P. Wadler. The call-by-need lambda calculus. *JFP*, 8(03), 1998.

John McCarthy. Recursive functions of symbolic expressions and their computation by machine, Part I. *CACM*, 3(4), 1960.

E. Meijer, B. Beckman, and G. Bierman. LINQ: reconciling object, relations and XML in the .NET framework. In *SIGMOD*. ACM, 2006.

Eugenio Moggi. Notions of computation and monads. *I&C*, 93(1), 1991.

Flemming Nielson and Hanne Riis Nielson. *Two-level functional languages*. Cambridge University Press, 2005.

John O'Donnell. Generating netlists from executable circuit specifications in a pure functional language. In *Functional Programming, Glasgow*. Springer, 1993.

A. Persson, E. Axelsson, and J. Svenningsson. Generic monadic constructs for embedded languages. In *IFL*, 2011.

F. Pfenning and C. Elliot. Higher-order abstract syntax. In *PLDI*, 1988.

Dag Prawitz. *Natural Deduction: A Proof-Theoretical Study*. Almqvist and Wiksell, Stockholm, 1965.

T. Rompf and M. Odersky. Lightweight modular staging: a pragmatic approach to runtime code generation and compiled DSLs. In *GPCE*. ACM, 2010.

Tiark Rompf. *Lightweight Modular Staging and Embedded Compilers: Abstraction without Regret for High-Level High-Performance Programming*. PhD thesis, EPFL, 2012.

Tiark Rompf, Nada Amin, Adriaan Moors, Philipp Haller, and Martin Odersky. Scala-virtualized: Linguistic reuse for deep embeddings. *HOSC*, 2013.

Arvind K Sujeeth et al. Composition and reuse with compiled domain-specific languages. In *ECOOP*. Springer, 2013.

J. Svenningsson and E. Axelsson. Combining deep and shallow embedding for EDSL. In *TFP*. Springer, 2012.

J. Svenningsson and B. Svensson. Simple and compositional reification of monadic embedded languages. In *ICFP*. ACM, 2013.

J. Svensson, M. Sheeran, and K. Claessen. Obsidian: A domain specific embedded language for parallel programming of graphics processors. In *IFL*. Springer, 2011.

Don Syme. Leveraging .NET meta-programming components from F#: integrated queries and interoperable heterogeneous execution. In *ML*. ACM, 2006.

Philip Wadler. Propositions as types. *CACM*, 2015. To appear.

Limsoon Wong. Normal forms and conservative extension properties for query languages over collection types. *JCSS*, 52(3), 1996.